

SPINES OF MINIMAL LENGTH

BRUNO MARTELLI, MATTEO NOVAGA, ALESSANDRA PLUDA, AND STEFANO RIOLO

ABSTRACT. In this paper we raise the question whether every closed Riemannian manifold has a spine of minimal area, and we answer it affirmatively in the surface case. On constant curvature surfaces we introduce the *spine systole*, a continuous real function on moduli space that measures the minimal length of a spine in each surface. We show that the spine systole is a proper function and has its global minima precisely on the extremal surfaces (those containing the biggest possible discs).

We also study *minimal spines*, which are critical points for the length functional. We completely classify minimal spines on flat tori, proving that the number of them is a proper function on moduli space. We also show that the number of minimal spines of uniformly bounded length is finite on hyperbolic surfaces.

1. INTRODUCTION

In this paper we study the spines of closed Riemannian manifolds that have minimal area with respect to the codimension-one Hausdorff measure. We describe the general setting, we prove that such spines exist on surfaces, and then we study the constant curvature case in detail.

In differential topology, a *spine* of a closed smooth n -manifold M is a smooth finite simplicial complex $P \subset M$ such that M minus a small open ball collapses onto P . In particular $M \setminus P$ is an open ball.

In all cases, we suppose that $\dim P < \dim M$, and this is the only restriction we make on dimensions: for instance, any point is a spine of the sphere S^n for all $n \geq 1$. The polyhedron P may also have strata of varying dimensions: for instance a natural spine of $S^2 \times S^1$ is the union of the sphere $S^2 \times q$ and the circle $p \times S^1$.

A compact manifold M has many different spines: as an example, one may give M a Riemannian structure and construct P as the cut locus of a point [8]. The manifold M has typically infinitely many pairwise non-homeomorphic spines, with portions of varying dimension.

The notion of spine is widely employed in topology: for instance, it may be used to define a *complexity* on manifolds [19, 20, 17], to study group actions [6] and properties of Riemannian manifolds [3]. In dimensions 2 and 3 spines (with generic singularities) arise naturally and frequently as the dual of 1-vertex triangulations. Topologists usually consider

2010 *Mathematics Subject Classification.* 57M50; 49Q20; 53A10.

Key words and phrases. spine; length functional; constant curvature surfaces.

spines only up to isotopy, and relate different spines (or triangulations) via some moves like “flips” on surfaces (see for instance [5, 11, 15, 23]) and Matveev-Piergallini moves [18, 21] on 3-manifolds.

However, it seems that spines have not been much studied from a geometric measure theory point of view, and this is the main purpose of this paper.

If M is a closed Riemannian manifold of dimension $n \geq 2$, every spine $P \subset M$ has a well-defined finite $(n-1)$ -dimensional Hausdorff measure $\mathcal{H}^{n-1}(P)$ called *area*. For instance, a point in S^n has zero area; the spine of $S^2 \times S^1$ described above has the same area 4π of S^2 .

We are interested here in the following problem.

Question 1.1. *Does every closed Riemannian manifold M of dimension $n \geq 2$ have a spine of minimal area?*

As an example, the answer is *yes* for all spheres S^n equipped with any Riemannian metric, since S^n has a spine of zero area (a point). The reader may notice here that we need to allow smaller-dimensional spines like points to get a positive answer to Question 1.1 on spheres (and also on other manifolds, see Section 2.1).

This paper is essentially devoted to surfaces: their spines have dimension ≤ 1 and it is of course more reasonable to employ the word *length* to indicate their area.

Theorem 1.2. *Every closed Riemannian surface S has a spine Γ of minimal length. The spine Γ is:*

- (1) *a point if S is diffeomorphic to a sphere,*
- (2) *a closed geodesic if S is diffeomorphic to a projective plane,*
- (3) *finitely many geodesic arcs meeting at trivalent points with angle $\frac{2}{3}\pi$ otherwise.*

The theorem says everything about spheres, so we restrict our attention to the other surfaces S . In analogy with minimal surfaces, we say that a spine $\Gamma \subset S$ is *minimal* if it is as prescribed in points (2) or (3) of the theorem: either a closed geodesic, or finitely many geodesic arcs meeting with angle $\frac{2}{3}\pi$ at trivalent points. These conditions are similar to that of having zero mean curvature for hypersurfaces. A spine of minimal length is minimal, but the converse may not hold: a minimal spine is a critical point of the length functional, while a spine of minimal length is a global minimum.

It is now natural to study these geometric objects on closed surfaces of constant curvature. Recall that all constant curvature metrics on a closed oriented surface S , considered up to orientation-preserving isometries and global rescalings, form the *moduli space* $\mathcal{M}(S)$ of S . The moduli space is not compact: on the torus T , the space $\mathcal{M}(T)$ is the $(2, 3, \infty)$ orbifold and is homeomorphic to a plane.

We completely classify all minimal spines in flat tori.




FIGURE 1. The number of minimal spines on each oriented flat torus in moduli space $\mathcal{M}(T)$, considered up to orientation-preserving isometries of T . The moduli space $\mathcal{M}(T)$ is drawn via the usual fundamental domain in \mathbb{H}^2 . At each point $z \in \mathcal{M}(T)$, the number $c(z)$ is the smallest among all numbers written on the adjacent strata (the function c is lower semi-continuous).

Theorem 1.3. *Every oriented flat torus T contains finitely many minimal spines up to orientation-preserving isometries. The number of such minimal spines is the proper function $c: \mathcal{M}(T) \rightarrow \mathbb{N}$ shown in Fig. 1.*

We see in particular that the square and the hexagonal torus at $z = i$ and $z = e^{\frac{\pi}{3}i}$ are the only flat tori that contain a unique minimal spine up to orientation-preserving isometries. We also discover that every flat torus T contains finitely many minimal spines up to isometry, but this number increases (maybe unexpectedly?) and tends to ∞ as the flat metric tends to infinity in moduli space.

By looking only at spines of minimal length we find the following.

Proposition 1.4. *Every oriented flat torus T has a unique spine of minimal length up to isometries of T . It also has a unique one up to orientation-preserving isometries, unless it is a rectangular non square torus and in this case it has two.*

Let $S: \mathcal{M}(T) \rightarrow \mathbb{R}$ be a function that assigns to a (unit-area) flat torus $z \in \mathcal{M}(T)$ the minimal length $S(z)$ of a spine in T . The function S may be called the *spine systole* because it is analogous to the (geodesic) systole that measures the length of the shortest

closed geodesic. We will prove that S has the following expression:

$$S(z) = \sqrt{\frac{1 + |z|^2 - |\Re(z)| + \sqrt{3}\Im(z)}{\Im(z)}}.$$

and has a unique global minimum at the hexagonal torus.

Now that the picture on the flat tori is perfectly understood, it is time to turn to higher genus hyperbolic surfaces and see if minimal spines have the same qualitative behavior there. Let S_g be a closed orientable surface of genus $g \geq 2$ and let $\mathcal{M}(S_g)$ be the moduli space of S_g . Every hyperbolic surface in $\mathcal{M}(S_g)$ has a spine of minimum length by Theorem 1.2 and its length defines a spine systole $S: \mathcal{M}(S_g) \rightarrow \mathbb{R}$.

An *extremal surface* is a hyperbolic surface that contains a disc of the maximum possible radius in genus g . Such surfaces were defined and studied by Bavard [7] and various other authors, see [13] and the references therein. We can prove the following.

Theorem 1.5. *The function $S: \mathcal{M}(S_g) \rightarrow \mathbb{R}$ is continuous and proper. Its global minima are precisely the extremal surfaces.*

A full classification of all minimal spines on all hyperbolic surfaces would of course be desirable; for the moment, we content ourselves with the following.

Proposition 1.6. *Every closed hyperbolic surface S has finitely many minimal spines of bounded length.*

In particular, we do not know if a closed hyperbolic surface has finitely many minimal spines overall (counting spines only up to isometry does not modify the problem, since the isometry group of S is finite).

In higher dimensions, Question 1.1 has already been addressed by Choe [9] who has provided a positive answer for all closed irreducible 3-manifolds. The techniques used in that paper are much more elaborate than the ones we use here.

2. PRELIMINARY DEFINITIONS

2.1. Spines. We recall some well-known notions that apply both in the piecewise-linear and in the smooth category of manifolds.

A *smooth finite simplicial complex* (or a *finite polyhedron*) in a smooth n -manifold M with (possibly empty) boundary is a subset $P \subset \text{int}(M)$ homeomorphic to a simplicial triangulation, such that every simplex is diffeomorphic through some chart to a standard simplex in \mathbb{R}^n .

The subset P has a well-defined *regular neighborhood*, unique up to isotopy: this is a piecewise-linear notion that also applies in the smooth category [14]. A regular neighborhood N of P is a compact smooth n -dimensional sub-manifold $N \subset M$ containing P in its interior which collapses simplicially onto P for some smooth triangulation of M ; as a consequence

$N \setminus P$ is an open collar of ∂N . If the manifold M itself is a regular neighborhood of P we say that P is a *spine* of M .

This definition of course may apply only when $\partial M \neq \emptyset$, so to define a spine P of a closed manifold M we priorly remove a small open ball from M . In particular we get that $M \setminus P$ is an open ball in this case.

In this paper we only consider spines P of closed manifolds M . We also consider by assumption only spines P with $\dim P < \dim M$, so that if M is endowed with a Riemannian metric P has a well-defined and finite $(n-1)$ -dimensional Hausdorff measure $\mathcal{H}^{n-1}(P)$ called *area*.

The area can be zero in some cases because we allow P to have dimension strictly smaller than $n-1$. For instance, any point is a spine of S^n and any complex hyperplane is a spine of complex projective space \mathbb{CP}^n , because their complements are open balls: in both cases the spine has zero area because it has codimension bigger than 1.

A spine may have strata of mixed dimensions: for instance a natural spine for $S^m \times S^n$ is the transverse union of two spheres $S^m \times q$ and $p \times S^n$ (the reader may verify that the complement is an open ball).

In the following we consider only the case in which the dimension of the manifold M is two.

Remark 2.1. The dimension of a spine Γ of a closed surface S is less or equal one and its homotopy type is completely determined by that of S . Indeed,

- Γ is homotopically equivalent to a point if and only if S is diffeomorphic to a sphere. Indeed, a regular neighbourhood of a point must be a disc.
- Γ is homotopically equivalent to a circle if and only if S is diffeomorphic to a projective plane. Indeed, a regular neighbourhood of a circle could be nothing but a Möbius strip, because an annulus has too many boundary components.
- Finally, assume that S is not diffeomorphic to a sphere, neither to a projective plane. A spine of S is an embedded graph $\Gamma \subset S$. Denoting with e the number of edges, with v the number of vertices of Γ and by χ the Euler characteristic, a necessary condition for Γ to be a spine of S is

$$\chi(S) - 1 = \chi(S - B^2) = \chi(\Gamma) = v - e. \quad (2.1)$$

2.2. Networks in Riemannian surfaces. Let S be a Riemannian surface. An embedded graph $\Gamma \subset S$ is a *network*: a union of a finite number of supports of simple smooth curves $\gamma^i : [0, 1] \rightarrow S$, intersecting only at their endpoints. A point in which two or more curves concur is called *multipoint*. Each curve γ^i of the network has length $\mathcal{H}^1(\text{Im}(\gamma^i))$ and the length of Γ is the sum of the lengths of all the curves, that is,

$$L(\Gamma) = \mathcal{H}^1\left(\bigcup_{i=1}^n (\text{Im}(\gamma^i))\right) = \sum_{i=1}^n \mathcal{H}^1(\text{Im}(\gamma^i)). \quad (2.2)$$

In the following, we will search the minima of this functional, restricted to the set of spines of a closed Riemannian surface S , endowed with the Hausdorff topology.

2.3. First variation of the length functional. Let Γ be a network in a closed Riemannian surface S . Let Φ_t with $t \in [0, T]$ be a smooth family of diffeomorphisms of S , with $\Phi_0(x) = x$ for all $x \in \Gamma$ and let $\Gamma_t = \Phi_t(\Gamma)$. Consider the vector field X on S defined as

$$X(x) = \frac{d}{dt} \Phi_t(x) \Big|_{t=0}.$$

The first variation formula for the length functional L for a network Γ is (see for instance [16]):

$$\frac{d}{dt} L(\Gamma_t) \Big|_{t=0} = \int_{\Gamma} \operatorname{div}_{\tau} X \, d\mathcal{H}^1, \quad (2.3)$$

where with $\operatorname{div}_{\tau}$ we denote the tangential divergence.

Let x_1, \dots, x_m be the multipoints of Γ . At every multipoint x_j , l_j curves concur. For $k = 1, \dots, l_j$ denote with τ_j^k the unit tangent vector to each of that curves at x_j . In particular,

$$\frac{d}{dt} L(\Gamma_t) \Big|_{t=0} = - \int_{\Gamma} H \cdot X \, d\mathcal{H}^1 + \sum_{j=1}^m \left(\sum_{k=1}^{l_j} \tau_j^k \cdot X \right), \quad (2.4)$$

where H is the curvature of the curve. A network Γ is *stationary* if there holds

$$\int_{\Gamma} \operatorname{div}_{\tau} X \, d\mathcal{H}^1 = 0 \quad \text{for every vector field } X \text{ on } S, \quad (2.5)$$

Thanks to (2.4), condition (2.5) is equivalent to require that each curve of the network Γ is a geodesic arc and the sum of the unit tangent vectors of concurring curves at a common endpoint is equal to zero.

We will actually consider only spines, and reserve the word *minimal* for a more restrictive configuration where only three edges concur at each multipoint, see Definition 3.2. The reason for that is that spines of minimal length will always be of this type.

2.4. Rectifiable sets and stationary varifolds. Let S be a Riemannian surface. We recall that a set $\Gamma \subset S$ is *countably 1-rectifiable* if it can be covered by countably many images of Lipschitz maps from \mathbb{R} to S , except a subset \mathcal{H}^1 -negligible.

Let Γ be a countably 1-rectifiable, \mathcal{H}^1 -measurable subset of S and let θ be a positive locally \mathcal{H}^1 -integrable function on Γ . Following [1, 2] we define the *rectifiable 1-varifold* (Γ, θ) to be the equivalence class of all pairs $(\tilde{\Gamma}, \tilde{\theta})$, where $\tilde{\Gamma}$ is countably 1-rectifiable with $\mathcal{H}^1((\Gamma \setminus \tilde{\Gamma}) \cup (\tilde{\Gamma} \setminus \Gamma)) = 0$ and where $\tilde{\theta} = \theta$ \mathcal{H}^1 -a.e. on $\Gamma \cap \tilde{\Gamma}$. The function θ is called *multiplicity* of the varifold.

A varifold (Γ, θ) is *stationary* if there holds

$$\int_{\Gamma} \theta \operatorname{div}_{\tau} X \, d\mathcal{H}^1 = 0, \quad (2.6)$$

for any vector field X on S . Notice that, if θ is constant and Γ is a network, condition (2.6) is consistent with (2.5). Recalling (2.3) this means that the network Γ is a critical point of the length functional.

3. RIEMANNIAN SURFACES

In this section we prove an existence result for the spines of minimal length of a closed Riemannian surface (Theorem 1.2).

Proposition 3.1. *Consider a closed Riemannian surface S not diffeomorphic to a sphere. Then there is a constant $K > 0$ such that $L(\Gamma) \geq K$ for every spine Γ of S .*

Proof. Let $r > 0$ be the injectivity radius of S . Every homotopically non-trivial closed curve in S has length at least $2r$. Every spine Γ contains at least one homotopically non-trivial embedded closed curve and hence has length at least $K = 2r$.

To verify the last fact, recall that $S \setminus \Gamma$ is an open 2-disc and hence the inclusion map $i: \Gamma \hookrightarrow S$ induces a surjection $i_*: \pi_1(\Gamma) \rightarrow \pi_1(S)$. Since $\pi_1(S) \neq \{e\}$ the spine Γ contains some homotopically non-trivial loop, and this easily implies that it also contains an embedded one. \square

Proof of Theorem 1.2. Thanks to the topological observations made in Subsection 2.1, the case of the sphere is trivial and that of the projective plane is well known. Therefore, we will henceforth suppose that S is neither diffeomorphic to a sphere nor to a projective plane. We want to prove that

$$\ell = \inf \{L(\Gamma) \mid \Gamma \text{ is a spine of } S\}. \quad (3.1)$$

is a minimum. We subdivide the proof into four steps.

Step 1: Γ_n *minimizing sequence of spines converges to a closed connected and rectifiable set Γ_∞ .*

We have just defined

$$\ell = \inf_{\Gamma \subset S} \mathcal{H}^1(\Gamma)$$

where Γ varies among all spines of S . Let Γ_n be a *minimizing sequence* of spines, that is a sequence such that $\mathcal{H}^1(\Gamma_n) \rightarrow \ell$. In particular, the sets Γ_n are closed, connected, rectifiable, and $S \setminus \Gamma_n$ is homeomorphic to an open disc for every n .

Thanks to Blaschke Theorem [4, Theorem 4.4.15], up to passing to a subsequence, the sequence converges $\Gamma_n \rightarrow \Gamma_\infty$ to a compact set Γ_∞ in the Hausdorff distance, and by Golab Theorem [4, Theorem 4.4.17] we get

$$\mathcal{H}^1(\Gamma_\infty) \leq \liminf_n \mathcal{H}^1(\Gamma_n) = \ell \quad (3.2)$$

and Γ_∞ is connected. Moreover by the Rectifiability Theorem [4, Theorems 4.4.7], the limit set Γ_∞ is rectifiable and connected by injective rectifiable curves.

Step 2: *Structure of Γ_∞ .*

We have proved that Γ_∞ has minimal length ℓ . We can associate to Γ_∞ a rectifiable varifold with multiplicity 1.

Since Γ_∞ is of minimal length, it follows that the corresponding varifold is stationary, hence is composed by a finite number of geodesic segments, joining finitely many nodes (see [2]). The minimality of Γ_∞ also implies that each node is the meeting point of exactly three curves, forming angles of $\frac{2\pi}{3}$ by standard arguments: suppose by contradiction that more than three curves concur at meeting point O . Consider a sufficiently small neighbourhood of this multiple point O where there are not other nodes. Take two segments concurring in O forming an angle $< \frac{2\pi}{3}$ and two points P and Q of these segments in the considered neighbourhood. Lift P and Q to the tangent plane of the surface at O through the exponential map. We know that in the tangent plane the Steiner configuration is the minimal length configuration joining O , P and Q . Replace the part of two curves on the surface joining P and Q with O in the neighbourhood with the image through the exponential map of the three segments of the Steiner configuration. Repeating iteratively this procedure, we obtain only triple junctions.

Step 3: *The set Γ_∞ intersects every homotopically non-trivial closed curve γ .*

If γ is a homotopically non-trivial closed curve, then each Γ_n intersects γ , because γ cannot be contained in the open disc $S \setminus \Gamma_n$. Hence, the Hausdorff limit Γ_∞ also intersects γ because γ is compact.

Step 4: *The set Γ_∞ is a spine.*

Let \mathcal{S} be the set of all closed subsets of Γ_∞ that intersect every homotopically non-trivial closed curve in S . By Zorn's lemma there is a $\Gamma \in \mathcal{S}$ which is minimal with respect to inclusion. We now prove that Γ is a spine.

The open set $S \setminus \Gamma$ contains no non-trivial closed curve. If $S \setminus \Gamma$ is an open disc we are done: we prove that this is the case. If one component U of $S \setminus \Gamma$ is not an open disc, it is not simply connected: hence it contains a simple closed curve γ which is homotopically non-trivial in U , but is necessarily trivial in S since it does not intersect Γ . Therefore γ bounds a disc $D \subset S$. Since $D \not\subset U$, the intersection $D \cap \Gamma$ is non-empty: if we remove $D \cap \Gamma$ from Γ we obtain another element of \mathcal{S} strictly contained in Γ , a contradiction.

Therefore $S \setminus \Gamma$ consists of open discs only. If they are at least two, there is an arc in Γ adjacent to two of them: by removing from Γ an open sub-arc in this arc we get again another element in \mathcal{S} strictly contained in Γ .

We know that $\Gamma \subset \Gamma_\infty$ is a spine, and hence $\Gamma = \Gamma_\infty$, for if not the length of Γ would be strictly smaller than that ℓ of Γ_∞ , a contradiction by (3.1). \square

Definition 3.2. A spine of a Riemannian surface S is *minimal* if it is a point, a closed geodesic, or if it is composed by finitely many geodesic arcs, meeting with angle $\frac{2\pi}{3}$ at trivalent points.

We have shown that a spine of minimal length is minimal. Of course, the converse may not hold. However, a minimal spine is a stationary point of the length functional thanks to (2.4).

Remark 3.3. Notice that the extension of the existence result in higher dimension present several difficulties: there is no higher dimensional version of Golab Theorem, because of the lack of semicontinuity of the Hausdorff measure. Also, it is not clear if a limit of spines is still a spine. However, as we already observed, the existence of a spine of minimal area in a closed irreducible 3-manifold has been proved in [9].

Remark 3.4. If the surface S is neither diffeomorphic to a sphere nor to a projective plane, we have shown that minimal spines are trivalent graphs. Hence, adding the equation $3v = 2e$ to (2.1), we get that the number of edges and that of vertices of a minimal spine are completely determined by the topology of S .

3.1. Non positive constant curvature surfaces. We restrict the attention to the case of non positive constant curvature surfaces. The goal is to show that minimal spines are local minimizers for the length functional, justifying our choice for the adjective “minimal”. Let us begin with a definition and a well-known lemma about the convexity of the distance function in \mathbb{H}^2 that we take from [10].

Definition 3.5. Let x_1, x_2 be points in $\mathbb{H}^2, \mathbb{R}^2$ or S^2 and $\lambda \in [0, 1]$. The *convex combination* $x = \lambda x_1 + (1 - \lambda)x_2$ is defined as follows:

$$\begin{aligned} \text{in } \mathbb{R}^2 : \quad & x = \lambda x_1 + (1 - \lambda)x_2 \\ \text{in } \mathbb{H}^2, S^2 : \quad & x = \frac{\lambda x_1 + (1 - \lambda)x_2}{\|\lambda x_1 + (1 - \lambda)x_2\|} \end{aligned}$$

where in the \mathbb{H}^2 case we are considering the hyperboloid model in \mathbb{R}^3 with the Lorentzian scalar product $\langle \cdot, \cdot \rangle$ and $\|v\| = \sqrt{-\langle v, v \rangle}$.

Lemma 3.6. Let x_1, x_2, y_1, y_2 be points of the hyperbolic plane \mathbb{H}^2 . For $\lambda \in (0, 1)$, consider the convex combinations $x_\lambda = \lambda x_1 + (1 - \lambda)x_2$ and $y_\lambda = \lambda y_1 + (1 - \lambda)y_2$. Then, we have

$$d(x_\lambda, y_\lambda) \leq \lambda d(x_1, y_1) + (1 - \lambda)d(x_2, y_2),$$

with equality only if x_1, x_2, y_1, y_2 belong to the same line. Here d denotes the distance in \mathbb{H}^2 .

Proof. Without loss of generality, for simplicity, we prove only the case $\lambda = \frac{1}{2}$, therefore x_λ (resp. y_λ) is the midpoint of x_1 and x_2 (resp. y_1 and y_2). If $x_\lambda = y_\lambda$ the theorem is trivial, hence we suppose $x_\lambda \neq y_\lambda$.

Let σ_p be the reflection at the point $p \in \mathbb{H}^2$. The map $\tau = \sigma_{y_\lambda} \circ \sigma_{x_\lambda}$ translates the line r containing the segment $x_\lambda y_\lambda$ by the quantity $2d(x_\lambda, y_\lambda)$: hence it is a hyperbolic transformation with axis r . We call $z_i = \tau(x_i)$ and note that $z_1 = \sigma_{y_\lambda}(x_2)$, hence $d(x_2, y_2) = d(z_1, y_1)$.

The triangular inequality implies that

$$d(x_1, z_1) \leq d(x_1, y_1) + d(y_1, z_1) = d(x_1, y_1) + d(x_2, y_2). \quad (3.3)$$

We notice that the equality holds only if x_1, y_1 and z_1 belong all to the same line.




FIGURE 2.

A hyperbolic transformation has minimum displacement on its axis r , hence

$$2d(x_\lambda, y_\lambda) = d(x_\lambda, z_\lambda) = d(x_\lambda, \tau(x_\lambda)) \leq d(x_1, \tau(x_1)) = d(x_1, z_1), \quad (3.4)$$

and the equality holds only if x_1 (and hence x_2) is in r . Finally we get $d(x_\lambda, y_\lambda) \leq \frac{1}{2}(d(x_1, y_1) + d(x_2, y_2))$ and hence d is convex.

Notice that the equality holds both in (3.3) and in (3.4) only if x_1, x_2, y_1, y_2 belong to the same line r . \square

Theorem 3.7. *Minimal spines of closed surfaces of non positive constant curvature are local minima for the length functional among spines, with respect to the Hausdorff distance.*

If the curvature is negative, these are strict local minima.

Proof. We prove the first statement by contradiction.

Consider a sequence of spines Γ_n of the surface S converging in the Hausdorff distance to a minimal spine Γ , such that $L(\Gamma_n) < L(\Gamma)$. The minimal spine Γ of S is a network composed by geodesic arcs joining k triple junctions x_1, \dots, x_k , where the number k depends only on the topology of S (see Remark 3.4). For n big enough, also Γ_n have k triple junctions $x_{1,n}, \dots, x_{k,n}$. Moreover, we can suppose that for n big enough Γ_n are composed only by geodesic segments. Indeed if Γ_n are not composed by geodesic segments, we can replace Γ_n with $\tilde{\Gamma}_n$, union of geodesic arcs, with the same triple junctions of Γ_n , and the value of the length functional decreases $L(\tilde{\Gamma}_n) \leq L(\Gamma_n) < L(\Gamma)$.

For n big enough, and for every $\lambda \in [0, 1]$ and $i \in \{1, \dots, k\}$, take the convex combination $x_{i,n}^\lambda = (1 - \lambda)x_i + \lambda x_{i,n}$ and define Γ_n^λ as the spine obtained by joining the points $x_{i,n}^\lambda$ with geodesic segments in the same pattern of Γ . We get a continuous family of spines $\{\Gamma_n^\lambda\}_{\lambda \in [0,1]}$ such that $\Gamma_n^1 = \Gamma_n$ and $\Gamma_n^0 = \Gamma$. By Lemma 3.6, the continuous function $F_n(\lambda) = L(\Gamma_n^\lambda)$ is convex (convexity of the distance function is easily proved also in the Euclidean case) and $F_n(1) \leq F_n(0)$. We also have $F'_n(0) = 0$ because the minimal spine Γ is a stationary point of the length functional. This implies that $L(\Gamma_n) = F_n(1) \geq F_n(0) = L(\Gamma)$ and we have a contradiction.

In the hyperbolic case, Lemma 3.6 provides strict convexity and hence Γ is a strict local minimum. \square

Remark 3.8. It is not restrictive to consider local minima of the length functional only among spines and not in the larger class of networks. Indeed, if we take a smooth family Φ_t of diffeomorphism of S with $t \in [0, T]$ and a spine Γ and we consider a small perturbation of Γ via these diffeomorphisms, $\Gamma_t = \Phi_t(\Gamma)$ is still a spine for t small enough. In particular, in the proof of Theorem 3.7, we show that $L(\Gamma) \leq L(\Phi_t(\Gamma))$, for all Φ_t and for t small enough.

4. FLAT TORI

In this section, we analyse minimal spines of the closed surface of genus 1: the torus $T = S^1 \times S^1$. In particular, we will fully determine all the minimal spines on T , endowed with any Euclidean metric.

4.1. Minimal spines of Riemannian tori. From Remark 3.4, for any Riemannian metric on the torus T , minimal spines have exactly 2 vertices and 3 edges. There are only two kinds of graph satisfying these properties: the θ -graph and the eyeglasses (see Figure 3).




FIGURE 3. The two trivalent graphs with three edges.

Both graphs can be embedded in a torus, but only the first in a way to obtain a spine. Indeed, it is easy to find a θ -spine in the torus (see Figure 4-a), and actually infinitely many non isotopic ones. Consider, instead, the eyeglasses as an abstract graph. Thickening the interiors of its edges, we get three bands. To have a spine on T , it remains to attach them, getting a surface homeomorphic to $T - B^2$. This is impossible. The band corresponding to one of the two “lenses” of the eyeglasses has to be glued to itself. There are only two ways to do this: one would give an unorientable surface (see Figure 4-b), the other would have too many boundary components (see Figure 4-c).

We have shown the following



FIGURE 4. a) A regular neighbourhood of a θ -spine on T . b)–c) An eyeglasses spine on T does not exist.

Proposition 4.1. *Trivalent spines on the torus are θ -graphs.*

Combined with Theorem 1.2, we get the following

Theorem 4.2. *On every Riemannian torus, minimal spines (exist and) are embedded θ -graphs with geodesic arcs, forming angles of $\frac{2\pi}{3}$ at their two meeting points.*

4.2. Basics on flat tori. From now, we consider only constant curvature Riemannian metrics on T . These correspond to flat (i.e. Euclidean) structures on T . We make a very rough introduction about the basic concepts of the topic to fix the notations, referring the reader to [10] for details.

Let us define

- the *Teichmüller space* \mathcal{T} of the torus as the space of isotopy classes of unit-area flat structures on T ,
- the *mapping class group* Mod of the torus as the group of isotopy classes of orientation preserving self-diffeomorphisms of T .
- the *moduli space* \mathcal{M} of the torus as the space of oriented isometry classes of unit-area flat structures on T .

The group Mod acts on the set \mathcal{T} and \mathcal{M} is the quotient by this action. The set \mathcal{T} is endowed with a natural metric (the Teichmüller metric), which makes it isometric to hyperbolic plane \mathbb{H}^2 . The action of the mapping class group can be identified with a discrete action by isometries.

Visualizing the hyperbolic plane with the upper-half space model $\{\Im(z) > 0\} \subset \mathbb{C}$, the action is given by integer Möbius transformations and the mapping class group is isomorphic to the group $SL_2(\mathbb{Z})$ of unit-determinant 2×2 integral matrices. The kernel of the action is $\{\pm I\}$, so that the moduli space \mathcal{M} has the structure of a hyperbolic orbifold, with $\pi_1 \mathcal{M} \simeq \text{PSL}_2(\mathbb{Z})$. A fundamental domain for the action is the hyperbolic semi-ideal triangle

$$D = \{|z| \geq 1, |\Re(z)| \leq 1\},$$

with angles $\frac{\pi}{3}$, $\frac{\pi}{3}$ and 0 (in grey in Figure 5). The quotient \mathcal{M} is the complete $(2, 3, \infty)$ -hyperbolic orbifold of finite area, with one cusp and two conical singularities of angles π and $\frac{2\pi}{3}$. We see \mathcal{M} as usual as D with the boundary curves appropriately identified.

We will also be interested in the following space:

- the *non-oriented moduli space* \mathcal{M}^{no} is the space of all isometry classes of (unoriented) unit-area flat structures on T .

We have $\mathcal{M}^{\text{no}} = \mathcal{M}/\iota$ where ι is the isometric involution that sends an oriented flat torus to the same torus with opposite orientation: the lift of ι to the fundamental domain D is the reflection with respect to the geodesic line $i\mathbb{R}^+$ and \mathcal{M}^{no} is the hyperbolic triangle orbifold

$$\mathcal{M}^{\text{no}} = D \cap \{\Re(z) \geq 0\}$$

with angles $\frac{\pi}{2}$, $\frac{\pi}{3}$ and 0. We do not need to fix an orientation on T to talk about spines, hence we will work mainly with \mathcal{M}^{no} .




FIGURE 5. Every $z \in \mathbb{H}^2$ in upper half-plane represents a flat torus obtained by identifying the opposite edges of the parallelogram with vertices $0, 1, z, z + 1$. A fundamental domain D for the action of $\mathrm{PSL}_2(\mathbb{Z})$ is colored in grey.

Every flat torus T has a continuum of isometries: the *translations* by any vector in \mathbb{R}^2 and the *reflections* with respect to any point $x \in T$. The tori lying in the mirror sides of $\mathcal{M}^{\mathrm{no}}$ have special names and enjoy some additional isometries:

- the *rectangular tori* are those lying in $i\mathbb{R}^+$,
- the *rhombic tori* are those in the other sides of $\mathcal{M}^{\mathrm{no}}$, namely $\{|z| = 1\} \cup \{\Re(z) = \frac{1}{2}\}$.

These flat tori are obtained by identifying the opposite sides of a rectangle and a rhombus, respectively. The tori in the cone points $z = i$ and $e^{\frac{\pi i}{3}}$ are the *square torus* and the *hexagonal torus*. On the hexagonal torus, the length d of the shortest diagonal of the rhombus equals the length l of any of its sides, while we have $d \geq l$ and $d \leq l$ on the sides $|z| = 1$ and $\Re = \frac{1}{2}$ respectively (we can call these rhombi *fat* and *thin*, respectively).

The rectangular and rhombic tori are precisely the flat tori that admit orientation-reversing isometries.

Teichmüller and moduli spaces have natural Thurston and Mumford-Deligne compactifications. In the torus case, these are obtained respectively by adding the circle “at infinity” $\partial\mathbb{H}^2 = \mathbb{R} \cup \{\infty\}$ to $\mathcal{T} = \mathbb{H}^2$ and a single point to \mathcal{M} or $\mathcal{M}^{\mathrm{no}}$. We denote the latter compactifications by $\overline{\mathcal{M}}$ and $\overline{\mathcal{M}^{\mathrm{no}}}$.

4.3. Hexagons. The study of minimal spines on a flat torus is intimately related to that of a particular class of Euclidean hexagons.

Definition 4.3. A *semi-regular hexagon* is a Euclidean hexagon with all internal angles $\frac{2\pi}{3}$ and with congruent opposite sides.

We define the *moduli space* \mathcal{H} as the space of all oriented semi-regular hexagons considered up to homotheties and orientation-preserving isometries. Similarly $\mathcal{H}^{\mathrm{no}}$ is defined by considering non-oriented hexagons and by quotienting by homotheties and all isometries. We get a map $\mathcal{H} \rightarrow \mathcal{H}^{\mathrm{no}}$ that is at most 2-to-1.

Two opposite sides of a semi-regular hexagon are parallel and congruent. A semi-regular hexagon is determined up to isometry by the lengths $a, b, c > 0$ of three successive sides, hence we get

$$\mathcal{H} = \{(a, b, c) \mid a, b, c > 0\} / \mathbb{R}_{>0} \times A_3$$

$$\mathcal{H}^{\text{no}} = \{(a, b, c) \mid a, b, c > 0\} / \mathbb{R}_{>0} \times S_3$$

where the multiplicative group of positive real numbers $\mathbb{R}_{>0}$ acts on the triples by rescaling, while A_3 and S_3 act by permuting the components.

We can visualize the space \mathcal{H}^{no} in the positive orthant of \mathbb{R}^3 by normalizing (a, b, c) such that $a + b + c = 1$ and $a \geq b \geq c$, and in this way \mathcal{H}^{no} is a triangle with one side removed (corresponding to $c = 0$) as in Fig. 6. The other two sides parametrize the hexagons with $a = b \geq c$ and $a \geq b = c$. The regular hexagon is of course $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. The space \mathcal{H}^{no} is naturally an orbifold with mirror boundary made by two half-lines and one corner reflector of angle $\frac{\pi}{3}$.

The space \mathcal{H} of oriented semi-regular hexagons is obtained by doubling the triangle \mathcal{H}^{no} along its two edges. Therefore \mathcal{H} is topologically an open disc, and can be seen as an orbifold where the regular hexagon is a cone point of angle $\frac{2\pi}{3}$. The map $\mathcal{H} \rightarrow \mathcal{H}^{\text{no}}$ may be interpreted as an orbifold cover of degree two.

The orbifold universal covering $\tilde{\mathcal{H}} \rightarrow \mathcal{H}$ is homeomorphic to the map $z \mapsto z^3$ from the complex plane to itself. The orbifold fundamental group $\pi_1 \mathcal{H}$ is isomorphic to the cyclic group $\mathbb{Z}/3$, acting on the plane by rotations of angle $\frac{2\pi}{3}$.




FIGURE 6. On the left, the orbifold structures on the moduli spaces of semi-regular hexagons, and their compactifications. On the right, a parametrization of \mathcal{H}^{no} by a triangle with one side removed in \mathbb{R}^3 .

Both \mathcal{H} and \mathcal{H}^{no} have natural compactifications, obtained by adding the side with $c = 0$, which consists of points $(a, 1 - a, 0)$ with $a \in [\frac{1}{2}, 1]$. Each such point corresponds to some “degenerate” hexagon: the points with $a < 1$ may be interpreted as parallelograms with angles $\frac{\pi}{3}$ and $\frac{2\pi}{3}$, while $(1, 0, 0)$ should be interpreted as a segment – a *doubly degenerate* hexagon. The underlying spaces of the resulting compactifications $\overline{\mathcal{H}}$ and $\overline{\mathcal{H}^{\text{no}}}$ are both homeomorphic to closed discs.

4.4. Spines. We now introduce two more moduli spaces \mathcal{S} and \mathcal{S}^{no} which will turn out to be isomorphic to \mathcal{H} and \mathcal{H}^{no} .

Let the *moduli space* \mathcal{S} be the set of all pairs (T, Γ) , where T is a flat oriented torus and $\Gamma \subset T$ a minimal spine, considered up to orientation-preserving isometries (that is, $(T, \Gamma) = (T', \Gamma')$ if there is an orientation-preserving isometry $\psi: T \rightarrow T'$ such that $\psi(\Gamma) = \Gamma'$). We define analogously \mathcal{S}^{no} as the set of all pairs (T, Γ) where T is unoriented, quotiented by all isometries. Again we get a degree two orbifold covering $\mathcal{S} \rightarrow \mathcal{S}^{\text{no}}$.

The opposite edges of a semi-regular hexagon H are congruent, and by identifying them we get a flat torus T . The boundary ∂H of the hexagon transforms into a minimal spine $\Gamma \subset T$ in the gluing process. This simple operation define two maps

$$\mathcal{H} \longrightarrow \mathcal{S}, \quad \mathcal{H}^{\text{no}} \longrightarrow \mathcal{S}^{\text{no}}.$$

Proposition 4.4. *Both maps are bijections.*

Proof. The inverse map is the following: given (T, Γ) , we cut T along Γ and get the original semi-regular hexagon H . \square

We will therefore henceforth identify these moduli spaces and use the symbols \mathcal{H} and \mathcal{H}^{no} to denote the moduli spaces of both semi-regular hexagons and pairs (T, Γ) . Of course our aim is to use the first (hexagons) to study the second (spines in flat tori).

4.5. The forgetful maps. The main object of this chapter is the characterization of the forgetful maps

$$p: \mathcal{S} \longrightarrow \mathcal{M}, \quad p: \mathcal{S}^{\text{no}} \longrightarrow \mathcal{M}^{\text{no}}$$

that send (T, Γ) to T forgetting the minimal spine Γ . The fiber $p^{-1}(T)$ over an oriented flat torus $T \in \mathcal{M}$ can be interpreted as the set of all minimal spines in T , considered up to orientation-preserving isometries of T . Likewise the fiber over an unoriented torus $T \in \mathcal{M}^{\text{no}}$ is the set of minimal spines in T , considered up to all isometries.

As we said above, we identify $\mathcal{H}, \mathcal{H}^{\text{no}}$ with $\mathcal{S}, \mathcal{S}^{\text{no}}$ and consider the compositions (which we still name by p)

$$p: \mathcal{H} \xrightarrow{\sim} \mathcal{S} \longrightarrow \mathcal{M}, \quad p: \mathcal{H}^{\text{no}} \xrightarrow{\sim} \mathcal{S}^{\text{no}} \longrightarrow \mathcal{M}^{\text{no}}.$$

The map p is described geometrically in Fig. 7.




FIGURE 7. The composition $p: \mathcal{H} \rightarrow \mathcal{S} \rightarrow \mathcal{M}$.




FIGURE 8. How to construct z from (a', b', c') : we pick the tripod with angles $\frac{2\pi}{3}$ and lengths a', b', c' , and place it in the upper half-plane so that the endpoints of the edges c' and b' lie in 0 and 1. The point z is the endpoint of the edge a' .

Proposition 4.5. *The map*

$$\begin{aligned} \tilde{p}: \mathcal{H}^{\text{no}} &\longrightarrow \mathbb{H}^2 \\ (a, b, c) &\longmapsto \frac{2c^2 - ab + ac + bc}{2(b^2 + c^2 + bc)} + i \frac{\sqrt{3}}{2} \frac{ab + bc + ac}{b^2 + c^2 + bc} \end{aligned}$$

is a lift of $p: \mathcal{H}^{\text{no}} \rightarrow \mathcal{M}^{\text{no}}$. The map \tilde{p} is injective and sends homeomorphically the triangle \mathcal{H}^{no} onto the light grey domain drawn in Fig. 9.

Proof. The map \tilde{p} is defined in Fig. 8. We first rescale the triple (a, b, c) , which by hypothesis satisfies $a \geq b \geq c$ and $a+b+c = 1$, to (a', b', c') , by multiplying each term by $1/\sqrt{b^2 + c^2 + bc}$. Now we pick the *tripod* with one vertex and three edges of length a', b', c' and with angles $\frac{2\pi}{3}$, and we place it in the half-space with two vertices in 0 and 1 as shown in the figure (we can do this thanks to the rescaling). The third vertex goes to some $z \in \mathbb{H}^2$ and we set $\tilde{p}(a, b, c) = z$.

The map \tilde{p} is clearly a lift of p and it only remains to determine an explicit expression for \tilde{p} . Applying repeatedly the Carnot Theorem, we get:

$$\begin{aligned}\operatorname{Arg}(z) &= \cos^{-1} \frac{2c^2 - ab + ac + bc}{2\sqrt{(b^2 + c^2 + bc)(a^2 + c^2 + ac)}}, \\ |z| &= \sqrt{\frac{a^2 + c^2 + ac}{b^2 + c^2 + bc}}.\end{aligned}$$

Finally,

$$\tilde{p}(a, b, c) = z = \frac{2c^2 - ab + ac + bc}{2(b^2 + c^2 + bc)} + i \frac{\sqrt{3}}{2} \frac{ab + bc + ac}{b^2 + c^2 + bc}.$$

The proof is complete. \square




FIGURE 9. The lift \tilde{p} sends \mathcal{H}^{no} homeomorphically to the light grey domain on the left. It sends the two sides $a \geq b = c$ and $a = b \geq c$ of \mathcal{H}^{no} to two geodesic lines in \mathbb{H}^2 , and sends the line at infinity $c = 0$ to the constant-curvature line dashed in the figure. In the oriented setting, the lift \tilde{p} sends \mathcal{H} to the whole bigger grey domain but is discontinuous on the segment s of hexagons of type $a = b \geq c$.

Fig. 9 shows that the lift \tilde{p} sends the two sides $a = b \geq c$ and $a \geq b = c$ of the triangle \mathcal{H}^{no} to geodesic arcs in \mathbb{H}^2 . Recall that the compactification $\overline{\mathcal{H}^{\text{no}}}$ is obtained by adding the singly-degenerate parallelograms $(a, 1 - a, 0)$ with $a \in [\frac{1}{2}, 1)$ and the doubly-degenerate $P = (1, 0, 0)$. The map \tilde{p} also extends to the parallelograms, and sends them to the dashed line in Fig. 8, but it does *not* extend continuously to P , not even as a map from $\overline{\mathcal{H}^{\text{no}}}$ to $\overline{\mathbb{H}^2}$. However, the map p from moduli spaces does extend.

Proposition 4.6. *The map $p: \mathcal{H}^{\text{no}} \rightarrow \mathcal{M}^{\text{no}}$ extends continuously to a map $p: \overline{\mathcal{H}^{\text{no}}} \rightarrow \overline{\mathcal{M}^{\text{no}}}$.*

Proof. Send the doubly degenerate point P to the point at infinity in $\overline{\mathcal{M}}$. \square

The oriented picture is easily deduced from the non-oriented one. We can lift the map $p: \mathcal{H} \rightarrow \mathcal{M}$ to a map $\tilde{p}: \mathcal{H} \rightarrow \mathbb{H}^2$ in Teichmüller space whose image is the bigger (both light and dark) grey domain in Fig. 9, however the map \tilde{p} is discontinuous at the segment consisting of all hexagons with $a = b \geq c$, which is sent to one of the two curved geodesic arcs in the picture. The map $p: \mathcal{H} \rightarrow \mathcal{M}$ extends continuously to a map $p: \overline{\mathcal{H}} \rightarrow \overline{\mathcal{M}}$.

The Möbius transformation $z \mapsto \frac{2iz + \sqrt{3} - i}{2z - 1 + \sqrt{3}i}$ gives a more symmetric picture of the image $\tilde{p}(\mathcal{H})$ inside the hyperbolic plane in the Poincaré disc model, as can be seen in Figure 10.




FIGURE 10. A picture of $\tilde{p}(\mathcal{H})$ in the Poincaré disc model.

Remark 4.7. Now, it is possible to quantify the “distance” between two spines on different flat tori, in a natural way. Indeed, thanks to the map p and the Teichmüller metric, the moduli space of spines can be endowed with a structure of hyperbolic orbifold.

For, identify – by a representation ρ – the action of the orbifold fundamental group $\pi_1\mathcal{H}$ on the universal covering $\tilde{\mathcal{H}}$ with the action on the hyperbolic plane of the group Γ , generated by an order-three elliptic rotation about the point $e^{i\pi/3}$. To get a developing map $d: \tilde{\mathcal{H}} \rightarrow \mathbb{H}^2$ for the hyperbolic structure on \mathcal{H} , lift ρ -equivariantly the map p .

The developed image $d(\tilde{\mathcal{H}})$ of \mathcal{H} is the interior in \mathbb{H}^2 of the regular non-geodesic ideal triangle joining the points ∞ , 0, 1 in Figure 10. The set $\tilde{p}(\mathcal{H})$ is a fundamental domain for the action of Γ and the space \mathcal{H} is identified with the quotient $d(\tilde{\mathcal{H}})/\Gamma$.

Hence, the moduli space of spines is an infinite-area incomplete hyperbolic orbifold. Its completion, supported on the complement of the doubly degenerate hexagon $\overline{\mathcal{H}} \setminus \{P\}$, is an orbifold with non-geodesic boundary. Its boundary is an infinite constant-curvature line whose points are (all the) four-valent geodesic spines with alternating angles $\frac{\pi}{3}$ and $\frac{2\pi}{3}$.

Similarly, $\overline{\mathcal{H}^{\text{no}}} \setminus \{P\}$ is a complete infinite-area hyperbolic orbifold with (unbounded, non-geodesic) boundary and two mirror edges.

4.6. The number of minimal spines. We are ready to determine the fiber $p^{-1}(T)$ for every flat torus T in \mathcal{M} and \mathcal{M}^{no} , which may be interpreted as the set of all minimal spines in T , considered up to (orientation-preservingly or all) isometries of T .




FIGURE 11. The number of minimal spines on each flat torus, in the unoriented and oriented setting. At each point z in \mathcal{M}^{no} or \mathcal{M} , the number $c^{\text{no}}(z)$ or $c(z)$ is the smallest number among all numbers written on the adjacent strata (both functions are lower semi-continuous).

Theorem 4.8. *Every unoriented flat torus T has a finite number $c^{\text{no}}(T)$ of minimal spines up to isometries of T . Analogously, every oriented flat torus T has a finite number $c(T)$ of minimal spines up to orientation-preserving isometries. The proper functions $c^{\text{no}}: \mathcal{M}^{\text{no}} \rightarrow \mathbb{N}$ and $c: \mathcal{M} \rightarrow \mathbb{N}$ are shown in Fig. 11.*

Proof. We can identify \mathcal{H}^{no} with its image $H = \tilde{p}(\mathcal{H}^{\text{no}})$ and note that the restriction of the projection $\pi: \mathbb{H}^2 \rightarrow \mathcal{M}^{\text{no}}$ to H is finite-to-one. Indeed, for $z, z' \in H$, $\pi(z) = \pi(z')$ if and only if $\Im z' = \Im z$ and $\Re z' = \pm \Re z + n$ for some integer n . The number $c^{\text{no}}(z)$ is the cardinality of the fiber $p^{-1}(z)$ and is easily shown to be as in Fig. 11-(left).

The oriented case is treated analogously: in that case $\tilde{p}(\mathcal{H})$ is the bigger grey zone in \mathbb{H}^2 and $p(z) = p(z')$ if and only if $z' = z + n$ for some $n \in \mathbb{Z}$. We get the picture in Fig. 11-(right). \square

The subtler cases are those where T has additional symmetries: in Fig. 12 we show the minimal spines of T as the metric varies from rectangular to square, fat rhombic, hexagonal, and finally thin rhombic.




FIGURE 12. The minimal spines in some oriented tori having additional symmetries. As we pass from rectangular to thin rhombic, we find 2, 1, 2, 1, and then 3 spines up to orientation-preserving isometries. They reduce to 1, 1, 2, 1, 2 up to all isometries, including orientation-reversing ones. Rectangles and rhombi that are thinner (ie longer) than the ones shown here have additional minimal spines that wind around the thin part.

4.7. Length. Let us define the *length* of a union of curves on a flat torus in the moduli space \mathcal{M} as its normalized length, that is its one dimensional Hausdorff measure on a unit-area representative torus. For example, if you consider the explicit representative given by the parallelogram \mathbb{R} -generated by 1 and $z \in \mathbb{C}$, you have to divide lengths by the square root of the area of that parallelogram, that is $\sqrt{\Im(z)}$.

Generalizing the well known concept for closed geodesics on surfaces (see [10]), we define the *length spectrum of minimal spines* of a flat torus T as the set of lengths of its minimal spines. Actually, we can define a more accurate spectrum, a set of triples corresponding to the lengths of the sides of any spine on T , in decreasing order, that is

$$\left\{ \frac{(a, b, c)}{\sqrt{(b^2 + c^2 + bc)\Im(\tilde{p}(a, b, c))}} \mid (a, b, c) \in p^{-1}(T) \right\}.$$

We have shown that both sets are finite, for any flat torus T .

In the following, we determine the spines of minimal length.

Theorem 4.9. *Every unoriented flat torus has a unique spine of minimal length up to isometry. In the oriented setting, instead, the same holds with the exception of the rectangular non square tori, for which there are exactly two.*

Proof. Fix a torus $T \in \mathcal{M}^{\text{no}}$. In the proof of Theorem 4.8, we observed that all points of $\pi^{-1}(T) \cap H \subset \mathbb{H}^2$ have the same imaginary part, so the lengths of the tripods we found with the map \tilde{p} can be compared without need of normalization. Clearly, only one is the shortest: that corresponding to $z_0 \in D \cap \{\Re(z) \geq 0\}$.

In the oriented case, consider the representative z_0 of $T \in \mathcal{M}$ with $z_0 \in D \setminus \{z \in \partial D \mid \Re(z) < 0\}$. There are three cases:

- if $\Re(z_0) > 0$, the unique shortest spine is that associated to z_0 itself;

- if $\Re(z_0) < 0$, it is associated to $z_0 + 1$;
- if $\Re(z_0) = 0$, that is T is a rectangular torus, both spines associated to z_0 and $z_0 + 1$ are of minimal length. These two spines are the same only for the square torus, indeed a rotation of $\frac{\pi}{2}$ sends one to the other.

These evident assertions can be verified by trigonometry, or by computing the lengths through the formula showed in the following . \square

The explicit expression on $\tilde{p}(\mathcal{H}) \subset \mathbb{H}^2$ for the *length function* $L: \mathcal{H} \rightarrow (0, +\infty)$, which to a spine assigns its length, can be simply computed by the law of sines:

$$L(z) = \sqrt{\frac{1 + |z|^2 - \Re(z) + \sqrt{3}\Im(z)}{\Im(z)}}.$$

The formula for the *spine systole* $S: \mathcal{M} \rightarrow (0, +\infty)$, a generalized systole which assigns to a flat torus the length of its shortest spines is

$$S(z) = \sqrt{\frac{1 + |z|^2 - |\Re(z)| + \sqrt{3}\Im(z)}{\Im(z)}}.$$

Both functions are proper and almost everywhere smooth. They extend to continuous functions $L: \overline{\mathcal{H}} \rightarrow (0, +\infty]$ and $S: \overline{\mathcal{M}} \rightarrow (0, +\infty]$.

For every continuous function $f: X \rightarrow \mathbb{R}$ on a topological 2-manifold X , recall that a point $p \in X$ is said to be *regular* if there is a (topological) chart around $p = 0$ in which $f(x, y) = \text{const} + x$. Otherwise, it is *critical*. A critical point $p \in X$ is said to be *non degenerate* if in a local (topological) chart around $p = 0$ we have either $f(x, y) = \text{const} - x^2 + y^2$, or $f(x, y) = \text{const} \pm (x^2 + y^2)$. The non degenerate critical points are necessarily isolated. The function f is *topologically Morse* if it is proper and all critical points are non degenerate. The classical Morse theory works also in the topological category.

Remark 4.10. The functions $L: \mathcal{H} \rightarrow (0, +\infty)$ and $S: \mathcal{M} \rightarrow (0, +\infty)$ are topologically Morse. For both functions, the set of sublevel k is empty if $k < \sqrt[4]{3}\sqrt{2}$, a point if $k = \sqrt[4]{3}\sqrt{2}$ and a topological disc otherwise. Hence, the unique critical point is a minimum. Therefore, among all minimal spines on all flat tori, exactly one is the shortest. As expected, it is the equilateral spine on the hexagonal torus. Its length is $\sqrt[4]{3}\sqrt{2} \approx 1.86$.

5. HYPERBOLIC SURFACES

Let now S_g be a closed orientable surface of genus $g \geq 2$. The oriented hyperbolic metrics on S_g form the *moduli space* $\mathcal{M}(S_g)$ and the minimum length of a spine furnishes the *spine systole*

$$S: \mathcal{M}(S_g) \longrightarrow \mathbb{R}.$$

We now prove some facts on the function S .

Theorem 5.1. *The function S is continuous and proper. Its global minima are precisely the extremal surfaces.*

Proof. The function S is clearly continuous because the length of spines varies continuously in the metric. We now prove properness as an easy consequence of the Collar Lemma [10].

By Mumford's compactness theorem the subset $\mathcal{M}_\varepsilon(S_g) \subset \mathcal{M}(S_g)$ of all hyperbolic metrics with (closed geodesic) systole $\geq \varepsilon$ is compact for all $\varepsilon > 0$, and the Collar Lemma says that for sufficiently small $\varepsilon > 0$ a hyperbolic surface $S \in \mathcal{M}(S_g) \setminus \mathcal{M}_\varepsilon(S_g)$ has a simple closed curve γ of length $< \varepsilon$ with a collar of diameter $C(\varepsilon)$, for some function C such that $C(\varepsilon) \rightarrow \infty$ as $\varepsilon \rightarrow 0$. Every spine Γ of S_g must intersect γ and cross the collar, hence $L(\Gamma) \geq C(\varepsilon)$ and therefore S is proper.

We now determine the global minima of S . Let Γ be a spine in $S \in \mathcal{M}(S_g)$ of minimal length. The spine has $6g - 3$ edges. As in the Euclidean case, by cutting S along Γ we get a hyperbolic polygon P with $12g - 6$ edges and all interior angles $\frac{2\pi}{3}$. The length of Γ is half the perimeter of P .

Porti has shown [22] that, among all hyperbolic n -gons with fixed interior angles, the one with smaller perimeter is the unique one that has an inscribed circle. Therefore among all polygons P with angles $\frac{2\pi}{3}$ the one that minimizes the perimeter is precisely the regular one R , that is the one whose sides all have the same length. We deduce that the global minima for S are the hyperbolic surfaces that have R as a fundamental domain, and these are precisely the extremal surfaces, as proved by Bavard [7]. \square

It would be interesting to investigate the function S and check for instance whether it is a topological Morse function, see [12].

In the flat case we have shown that the number of minimal spines is finite. In the hyperbolic setting, we do not know if the same is true. To conclude the section, we prove a partial result:

Theorem 5.2. *The number of minimal spines with bounded length of a closed hyperbolic surface S is finite.*

Proof. In the proof of Theorem 1.2, it results clear that every set of minimal spines of equibounded length is compact. We now prove that for hyperbolic surfaces every such set is discrete, hence finite. By contradiction, let Γ_n be a sequence of distinct minimal spines of S of equibounded length, converging in the Hausdorff distance to the minimal spine Γ . Moreover $L(\Gamma_n) \rightarrow L(\Gamma)$. For every $\lambda \in [0, 1]$ and every n big enough, we construct, exactly as in the proof of Theorem 3.7, the (not necessarily minimal) spine Γ_n^λ and continuous function $F_n(\lambda) = L(\Gamma_n^\lambda)$. The surface S is hyperbolic, therefore, by Lemma 3.6, $F_n(\lambda)$ is strictly convex. Both Γ and Γ_n are minimal spines, that is stationary points of the length functional, hence, $F_n'(0) = F_n'(1) = 0$ and F_n is constant in λ : a contradiction. \square

REFERENCES

- [1] W. K. ALLARD, *On the first variation of a varifold*, Annals of Math. **95** (1972), 417–491.

- [2] W. K. ALLARD – F. J. ALMGREN, *The Structure of Stationary One Dimensional Varifolds*. Invent. Math. **34** (1976), 83–97.
- [3] S. B. ALEXANDER – R. L. BISHOP, *Spines and Homology of Thin Riemannian Manifolds with Boundary*, Adv. Math. **155** (2000), 23–48.
- [4] L. AMBROSIO – P. TILLI. “Topics on Analysis in Metric Spaces,” Oxford University Press, 2004.
- [5] J. ARAMAYONA – T. KOBERDA – H. PARLIER, *Injective maps between flip graphs*, [arXiv:1409.7046](#)
- [6] A. ASH, *Small-dimensional classifying spaces for arithmetic subgroups of general linear groups*, Duke Math. J. **51** (1984), 459–468.
- [7] C. BAVARD, *Disques extrémaux et surfaces modulaires*, Ann. Fac. Sci. Toulouse Math. **2** (1996), 191–202.
- [8] M. A. BUCHNER, *Simplicial structure of the real analytic cut locus*, Proc. Amer. Math. Soc. **64** (1977), 118–121.
- [9] J. CHOE, *On the existence and regularity of fundamental domains with least boundary area*, J. Differential Geom. **29** (1989), 623–663.
- [10] B. FARB AND D. MARGALIT “A primer on mapping class groups,” Princeton University Press, Princeton, NJ **49** (2012).
- [11] S. FOMIN – M. SHAPIRO – D. THURSTON, *Cluster algebras and triangulated surfaces. I. Cluster complexes*, Acta Math. **201** (2008), 83–146.
- [12] M. GENDULPHE, *The injectivity radius of hyperbolic surfaces and some Morse functions over moduli spaces*, [arXiv:1510.02581](#)
- [13] E. GIRONDO – G. GONZALEZ-DIEZ, *On extremal Riemann surfaces and their uniformizing Fuchsian groups*, Glasgow Math. J. **44** (2002), 149–157.
- [14] M. H. HIRSCH, *Smooth regular neighborhoods*, Ann. Math. **76** (1962), 524–530.
- [15] M. KORKMAZ – A. PAPADOPOULOS, *On the ideal triangulation graph of a punctured surface*, Ann. Inst. Fourier **62** (2012), 1367–1382.
- [16] F. MAGGI *Sets of Finite Perimeter and Geometric Variational Problems. An introduction to geometric measure theory*. Cambridge University Press, Cambridge (2012).
- [17] B. MARTELLI, *Complexity of PL manifolds*, Alg. & Geom. Top. **10** (2010), 1107–1164.
- [18] S. MATVEEV, *Transformations of special spines and the Zeeman conjecture*, Math. USSR-Izv. **31** (1988), 423–434.
- [19] ———, *Complexity theory of three-dimensional manifolds*, Acta Appl. Math. **19** (1990), 101–130.
- [20] ———, “Algorithmic Topology and Classification of 3-Manifolds”, second edition, Algorithms and Computation in Math. **9**, Springer, Berlin (2007).
- [21] R. PIERGALLINI, *Standard moves for standard polyhedra and spines*, Rendiconti Circ. Mat. Palermo **37** (1988), 391–414.
- [22] J. PORTI, *Hyperbolic polygons of minimal perimeter with given angles*, Geom. Dedicata **156** (2012), 165–170.
- [23] D. SLEATOR – D. D. TARJAN – W. P. THURSTON, *Rotation distance, triangulations, and hyperbolic geometry*, J. Amer. Math. Soc. **1** (1988), 647–681.

DIPARTIMENTO DI MATEMATICA “TONELLI”, LARGO PONTECORVO 5, 56127 PISA, ITALY
E-mail address: martelli at dm dot unipi dot it

DIPARTIMENTO DI MATEMATICA “TONELLI”, LARGO PONTECORVO 5, 56127 PISA, ITALY
E-mail address: novaga at dm dot unipi dot it

DIPARTIMENTO DI MATEMATICA “TONELLI”, LARGO PONTECORVO 5, 56127 PISA, ITALY
E-mail address: pluda at mail dot dm dot unipi dot it

DIPARTIMENTO DI MATEMATICA “TONELLI”, LARGO PONTECORVO 5, 56127 PISA, ITALY
E-mail address: riolo at mail dot dm dot unipi dot it